

Final Progress Report

Design Optimization of a GaAs RISC Microprocessor with Area-Interconnect MCM Packaging

Grant Number: DAAH04-94-G-0327

Richard B. Brown

University of Michigan

Foreword

The original objectives of this project were to develop the technologies and design automation environment for high clock-rate MCM-packaged gallium arsenide circuits which used flip-chip array I/O interconnect, and to demonstrate these technologies in the a prototype microprocessor. CAD tools were to be developed to support optimization of such systems for performance, power and cost.

The project involved close collaboration with two subcontractors, Motorola Semiconductor for the complementary gallium arsenide (CGaAs) technology, and Cascade Design Automation for CAD tools. Motorola scaled the CGaAs process from 0.7 μm minimum dimensions to 0.5 μm , improved the CGaAs yield significantly, conducted experiments in reducing threshold voltages, and fabricated prototype circuits. In the early days of the project, Motorola's Space and Systems Technology Group was a regular subcontractor, but when the budget was reduced, they agreed to continue to collaborate with us, bearing all of the costs themselves; most of the work was done under this arrangement. Cascade provided the physical design tools and support for these tools throughout the project. In addition, they developed a special tool for placing arrayed I/O pads on chips for flip-chip assembly. At the University of Michigan, a PowerPC architectural simulator was developed to evaluate cycles-per instruction for various microarchitectures, a CGaAs cell library was developed, and all of the design and testing of circuits was done.

Early in the project, the goal was to design the prototype system to operate with a 1 GHz clock, and advances were made in complementary GaAs technology, circuits, and packaging to enable this. When Motorola decided that CGaAs was the technology for their proposed Celestri satellite system, they fixed the minimum dimensions at 0.5 μm , rather than further scaling the process, and froze the thresholds at ± 0.55 V, rather than reducing them; on the positive side, they added a low-temperature GaAs buffer layer under the devices, which improved subthreshold slope and made the transistors extremely radiation hard to single-event upset. These decisions eliminated the possibility of building a fast processor in CGaAs, so at this point, the focus of the prototype system was shifted to space applications, which could take advantage of both the radiation hardness, and the excellent power-delay product of CGaAs.

CGaAs technology was analyzed to determine the most cost-effective scaling factor for each design rule; the methodology and tools developed for this can be applied also to the nonlinear scaling of deep-submicron CMOS processes. Static, domino and dual-rail domino (CVSL) circuits were designed to evaluate CGaAs for use in VLSI digital circuits. Phase-Locked Loop and current-mode I/O circuits were designed and tested. To facilitate the design of high-performance integrated circuits, logic synthesis and place and route tools were written. A gold-bumping process was developed in the UM solid-state electronics laboratory which produces bumps on pitches as tight as 50 μm . A superscalar PowerPC microarchitecture was developed for implementation in CGaAs with its modest integration levels, and fabricated in CMOS to prove correctness of the design. The project culminated in the design and testing of the radiation-hard CGaAs PUMA PowerPC microprocessor, which incorporates an area-I/O array.

DTIC QUALITY INSPECTED 4

20000628 067

Final Progress Report

Table of Contents

Foreword	1
Statement of Problem	3
Summary of Results	4
List of Publications	10
List of Participants	15
Report of Inventions	17
Bibliography	17

Final Progress Report

Statement of Problem

Microprocessors have had a profound impact on both the scientific world and on our personal lives. Through astonishing advances in performance, they have replaced traditional mainframes and supercomputers with microprocessor-based workstations and servers [1]. The remarkable decrease in cost vs. performance for microprocessors has made computing ubiquitous in our society.

Processor trends can be seen by surveying the microprocessors presented at ISSCC. Over the ten years before this project began, gate delays improved at 12% per year; clock frequencies increased at 40% per year; and transistor-counts grew at 40% per year [2]. The performance of systems made from these processors, as measured by the integer SPEC benchmarks, improved at a compounded rate of 59% per year, resulting in an increase in computing power of more than 100 fold over the decade.

The disparity in improvement between gate delay and clock frequency was due to the fact that some of the additional transistors made available during these years were used to pipeline processors, reducing the number of gates between latches. However, with processors having on the order of ten pipeline stages, additional pipeline depth provides diminishing performance returns, and could not be expected to maintain the steep increase in clock frequency seen before. The growing transistor budget also supported the addition of on-chip cache memory, which reduced load/store latency to memory. But again, with large first- and second-level instruction and data caches on chip, the performance return for enlarging cache memories beyond their current sizes was modest. To keep new processors on the performance curve, architects also invested their additional transistors in multiple functional units for concurrent instruction execution (superscalar architectures). Unfortunately, the benefits of parallelism also diminish with scale in general purpose machines.

With pipelining, on-chip cache size, and instruction issue width all at their points of diminishing return, semiconductor manufacturers turned to technology to keep the growth in computer performance on the curve. Increased attention to the scaling of CMOS, the inclusion of more fine-pitch metal interconnect layers, and more aggressive circuit techniques allowed vendors to increase the clock frequency and thereby increase processor throughput. The importance of semiconductor technology to future high-end computer performance warranted the evaluation of processes such as Complementary GaAs and SOI, which were outside of the mainstream. CGaAs has the switching speed of HEMTs with many of the circuit advantages of CMOS. The low-voltage operation of CGaAs, combined with the good switching speed, give CGaAs an excellent power-delay product. Process changes during this project made it extremely radiation hard.

When this project began, CGaAs had been targeted primarily at RF applications, with little digital work having been done. This project aimed to explore CGaAs technology for VLSI digital circuits, and to provide the packaging technologies needed to make it viable in such applications.

Final Progress Report

Summary of Results

This section of the report includes an overview of CGaAs technology, a list of accomplishments, and conclusions that can be drawn from the work in this project.

Overview of CGaAs Technology

CGaAs, a complementary heterostructure-insulated-gate FET technology, has been described in some detail in [3–5]. A sketch of the device structure is shown in Fig. 1. CGaAs integrates an enhancement-mode P-channel HFET with a high performance N-channel HFET. Historically, the primary interest in GaAs and other III/V materials has derived from their high electron mobilities. While holes in III/V materials do not enjoy an intrinsic mobility advantage over those in silicon, the pseudomorphic P-channel HFETs in this process have three to five times higher transconductance at given gate dimensions than their silicon counterparts. As seen in Fig. 1, the CGaAs process uses epitaxially-grown wafers, the cost of which includes both that of the initial GaAs wafer and of growing the additional layers. In moderate volumes, this can be 20 to 25 times the cost of a silicon substrate. Though the wafers are smaller and more expensive than silicon, the CGaAs process requires only 13 masks through three levels of interconnect, compensating in part with process efficiency for the more costly starting material. CGaAs, however, has not enjoyed the efficiencies of high-volume production as has CMOS, and completed wafers from a high volume CMOS process line cost approximately 40% to 50% as much as completed CGaAs wafers. Considering all of these factors, the price of a finished complementary gallium arsenide die is approximately 4.8 times the price of a similar-size, high volume CMOS die.

Standard gate lengths were scaled from 0.7 to 0.5 μm , and experimental N-channel devices at 0.35 μm gate lengths performed well. CGaAs has a power-delay product of 0.01 mW/MHz/gate at 0.7- μm minimum feature size. Recent improvements to the epi structure make the CGaAs process resistant to single-event upset (fewer than 10^{-9} Upsets/Bit-Day for SCFL logic and 10^{-10} Upsets/Bit-Day for complementary logic), as well as to large total dose radiation (more than 10^8 Rads) and latchup (more than 10^{12} Rads). Typical parameters for 0.7- μm channel-length devices having ± 0.55 V thresholds (measured with $V_{\text{dd}} = 1.5$ V) are given in Table 1. As seen in the table, both N and P channel devices have good output conductances and pinch-off characteristics. The original device thresholds of ± 0.55 V were selected because they yielded the optimum power-delay product in complementary circuits; they produce a drain-current ratio between N and P devices of about 4:1.




Fig. 1: CGaAs process cross-section.

Final Progress Report

Table 1: CGaAs Device Parameters.

Parameter	NFET (0.7x10 mm)	PFET (0.7x10 mm)
V_{th} (V)	+0.55	-0.55
I_{dss} (mA)	1.8	0.5
g_m (mS/mm)	280	60
Beta (mA/V ² -mm)	270	50
Subth slope (mV/dec)	75	90
Subth Current (nA) ($V_{gs}=0V$)	< 1	< 10

The low threshold voltages and high transconductance of CGaAs result in good performance at low voltages. Fig. 2 shows unloaded ring oscillator delays versus supply voltage for several logic families (thresholds of +/- 0.55V). The delay of 1.0- μ m CGaAs is less than that of 0.5- μ m CMOS or Thin-Film Silicon-on-Insulator (TFSOI), and the 0.5- μ m CGaAs shows delays below 100 ps with a 1.2V power supply. Power dissipation is not indicated in the figure. Lower threshold voltages will make the CGaAs circuits faster yet.

Two key parameters of concern in complementary heterostructure FET devices are gate leakage and sub-threshold drain-source leakage [3, 6], which determine the stand-by power dissipation of complementary circuits. Unlike Si CMOS, which has an SiO₂ gate insulator, the CGaAs gate is a Schottky diode to AlGaAs. Substantial gate current flows for gate voltages in excess of about one volt. Gate leakage current depends on the Schottky barrier height and band offsets. The large valence band offset (about 0.55 V) of high mole-fraction AlGaAs, as used in these devices, improves the gate leakage of PFETs. Typically, the PFET gate-diode turn-on voltage, defined as the gate voltage resulting in 1 μ A/ μ m² gate area at $V_{ds} = 0$, is -2V. NFETs have a turn-on voltage of 1.75 V. The gate turn-on voltages are also influenced by implant straggle effects. Drain-induced barrier lowering increases gate current when the drain-to-source voltage is high (which occurs when a logic input changes state).




Fig. 2: Propagation delay of CGaAs, CMOS, and TFSOI versus Supply Voltage. Gates are driving one load.

Final Progress Report

Accomplishments

The power efficiency and radiation hardness of CGaAs make it attractive for space and satellite applications. However, CGaAs presents design challenges such as reduced power-supply voltage (little headroom in circuits), proportionately large threshold voltages (lower speed than could otherwise be achieved by the HEMTs), gate and subthreshold drain-source leakage (higher power), and low integration levels (restrictions on architectures). CGaAs technology has been studied in this project for implementing large VLSI circuits such as microprocessors, in light of these challenges.

The Motorola 0.5- μm CGaAs process, which had been developed by shrinking the gate length of a 0.7- μm process, was the primary technology employed in this project. It was clear from the beginning that the design rules were not optimal and that the process needed to be scaled. A considerable amount of process development work was done on CGaAs to shift the thresholds, before the decision was made that they would be fixed at ± 0.55 V in order to assure that circuits could be delivered on schedule for the Celestri program. The yield was significantly improved on the CGaAs process through a defectivity program, and subthreshold leakage was reduced by adding a low-temperature GaAs buffer layer. A negative impact of this layer was that it reduced transistor gain, but because it is characterized by a very short carrier lifetime, it provides single-event upset protection for the process, which, because it has no SiO_2 gate oxide nor device isolation, has always been intrinsically hard to total radiation dose effects.

As CMOS processes are shrunk below 0.18 μm , the linear scaling of some design rules will be very difficult, so non-linear scaling will be needed for CMOS in the near future. Working with Motorola process engineers, we evaluated CGaAs for scaling. In doing so, we developed a general (works for any technology) methodology for quantitatively evaluating semiconductor processes for optimal scaling. The methodology includes identifying the design rules which have the greatest impact on the scaling objective and analyzing the area, power and performance improvements as these rules are incrementally scaled. The improvement data is combined with die cost estimates to produce a cost/benefit ratio which can guide scaling decisions. The methodology is based on the automated analysis of embedded static RAMs generated by a process-independent, optimizing SRAM compiler developed as part of this project. A cost/benefit analysis of the CGaAs design rules shows that when operating under a fixed spending cap, this nonlinear scaling approach can provide greater improvements in area and performance than linear scaling. The analysis results for the 0.5- μm CGaAs process recommend that threshold voltages be reduced, and that the first of a number of recommended scaling steps should be a 30% reduction of the source drain area and via/metal pitch.

Full complementary, unipolar (pseudo direct-coupled FET logic), pass-gate logic, and domino logic styles were evaluated in the complementary GaAs technology. A logic-evaluation test chip was fabricated at Motorola. Because initial evaluations of dynamic logic yielded promising results, a PowerPC ALU was designed in Domino logic. While this circuit was in fabrication, a yield-compromising design rule problem was identified; it became necessary to break gate-metal runs between n- and p-transistors to avoid leakage paths. This test run did provide valuable experience with the various logic styles, but because of the design-rule problem, did not yield on the dynamic ALU. An environment to help circuit designers optimize transistor sizes in SPICE netlists over power, area and delay was developed as part of this effort. Changes in focus at Motorola during this time caused the project to shift from high-clock rate designs to radiation-hard applications.

The original plan included implementation of a virtual memory system, and a software-managed in-cache translation mechanism for the processor was developed. This is an extremely low-overhead memory management scheme which provides all the benefits of traditional schemes but removes a substantial amount of hardware from the critical path, enabling much faster clock speeds. When the project budget was reduced with a reorganization at DARPA, we dropped the virtual memory and floating-point unit.

A trace-driven architectural simulator was developed to guide the design. To verify functionality of the PUMA design, we developed a random instruction generator which produces code based on a user-specified maximum number of loops and branches, and on flags specifying whether to use unimplemented instructions and misaligned memory accesses. Certain classes of instructions can be exercised, and register usage can be limited in order to stress forwarding interlocks. Simulation results with this code running on the Verilog PUMA model and on a PowerPC architectural simulator are compared to verify proper functionality.

Final Progress Report

Architectural methods of enhancing processor performance within the constraint of limited on-chip cache were explored. A method of prefetching called 'runahead' allows the processor to execute instructions under a cache miss, exposing other loads and stores that might have also generated cache misses, so that these can be prefetched. A second approach we have evaluated scans the instruction stream for branches as the instruction cache is loaded, and uses branch-prediction information to prefetch further instructions.

Development of the PUMA processor architecture was driven by the limited CGaAs integration level. The processor is implemented with a small on-chip primary instruction cache and a larger off-chip primary data cache. The instruction fetch mechanism is guided by an efficient two-level dynamic branch predictor and branch target buffer. Computation is performed by a small superscalar execution core comprised of branch, arithmetic, and load/store units. Based on trace-driven simulations of standard benchmark programs, the architecture should achieve 0.77 instructions per cycle. Out-of-order execution is supported by dedicated reservation stations for each functional unit and an eight-entry reorder buffer. The decode process translates complex PowerPC instructions into one or more simple RISC operations. A 0.35 μm CMOS version of the architecture was first prototyped. It has 280 pins, measures 9.9 x 9.9 mm, and contains 830K transistors. The chip was packaged in a 391-pin ceramic PGA. The chip is not fully tested yet, but so far, no errors have been detected.

The project culminated in the design and testing of the radiation-hard CGaAs PUMA PowerPC microprocessor, shown in Fig. 3, which incorporates an area-I/O array. This CGaAs version was further simplified to meet an integration limit of 400,000-transistors: the data cache was moved off chip, out-of-order execution was eliminated, and the architecture was modified to be single-issue. The CGaAs chips were fabricated at Motorola and ten of these chips were assembled (using just the peripheral I/O) in conventional PGA packages for initial testing. Of the ten microprocessors packaged, all of them passed basic tests but only two sequenced and executed instructions properly.




Fig. 3: CGaAs Power PC.

Final Progress Report

The two chips that passed had varying degrees of success with more advanced tests. None of the devices passed all the tests completely. Immediate instructions worked and program address sequencing worked, but instructions that manipulate register data gave bad data out. Functionality of the ALU, load/store unit, and the branch unit can be inferred from these tests, however; since output data is often bad it is not known if errors are introduced by registers or the buses. The branch instructions did work successfully. Using branches, the critical path of the FXU could be tested. The FXU operates at maximum frequencies of 42 MHz on chip one and 33 MHz on chip two. This test exercises only the critical path in the branch unit with certainty. There is not much difference in power dissipation between operating frequencies, indicating that most of the power is dissipated as static power. Eighteen percent of the power is dissipated in the core, the remainder in the pads. At a nominal operating voltage of 1.3 V, the FXU can be run optimally at 20 MHz dissipating 274 mW.

None of the devices passed the instruction cache tests indicating non-functional caches. More detailed cache testing was performed on a separate 2 Kbyte SRAM chip. It used the same SRAM design as the FXU caches. These chips also failed. The data out always followed the data in, indicating that the decoder was not working correctly. The decoder uses DCFL NOR gates. The ratios of these gates were not sufficient to provide a low enough output low voltage over process corners. Process data showed that the beta values, drive currents, and leakage currents of the N and P transistors as well as the threshold voltage of the P devices had a much wider distribution than anticipated. The degradation of the P device indicated by process data could also explain the other results from the testing. Leakage currents would be higher and some gates may not turn off at all, adding to static power and data errors. Further testing of the scan path and circuit simulations with the measured process corners should help identify the exact problems. Unfortunately, with the collapse of the Celestri project, Motorola is no longer running the CGaAs process, so there is no opportunity to modify the circuits for another run, and no chance of getting tighter process parameter control.

The PUMA project has developed new packaging and I/O signalling capabilities which are appropriate for military and aerospace applications now, as well as for future commercial CMOS systems. The processor chip includes a 315-pin area I/O pad array with a pad pitch of 6 mils, in addition to 288 pads in a staggered peripheral ring. It is designed for flip-chip assembly using gold bumps on a fine-pitch MCM-L board, connecting it to level-1 data cache, a memory management unit, PCI interface, and unified level-2 cache. The gold bumping process, which makes precisely-sized bumps of the desired aspect ratios, was developed in the University of Michigan Solid-State Electronics Laboratory. A multichip module, fabricated by Micromodule Systems, is a test vehicle for exploring design issues such as flip-chip area array attachment for more than 1,000 pads, minimum feasible pad pitch, and pad yields for various pad pitches (50, 75, 100, 125, and 175 μm pitch).

The PUMA project has also contributed to high-performance signalling technology. CGaAs Gunn transceiver, differential voltage, and switched current I/O interfaces have been designed, fabricated and tested. Test results indicate that these circuits in CGaAs can support bit rates of at least 650 Mb/s/pin (limited by the test set-up). An advanced transceiver based on switched-current techniques has also been designed. The receiver actively terminates the input line to its characteristic impedance using an active current mirror. The transmitted current pulse is 1.5 mA. The receiver is biased using a feedback circuit that overcomes parametric variations between the transmitting and receiving chips; it compensates for processing variations by adjusting the bias levels of the receiving chip. Simulations indicate that the circuit can support 1.2 Gb/s/pin signaling while dissipating only 3.3 mW, with a 1.4 V supply. A CGaAs delay-locked-loop (DLL) has been designed to explore the effects of low supply voltage and headroom on phase noise performance. Simulations indicate that the DLL would operate at 500 MHz, with a peak jitter of 88pS.

A CGaAs PLL clock generator was also designed and tested in this project. It operated at up to 800 MHz with a 1.5 V supply and 120 ps phase jitter. The CGaAs design was operational at a supply voltage as low as 0.8 V. A test MCM and CGaAs driver and receiver chips were designed for use with this PLL, to evaluate MCM signal integrity with low-voltage, high-edge-rate signals, and to test various driver and receiver circuits. The MCM was fabricated at MicroModule Systems (MMS), through Midas. The MCM included Mayo-designed passive test structures for measuring the MCM interconnect properties. The PLL clock generator was designed to phase lock to a low-speed input clock and produce a programmable multiple of this frequency for use as the GaAs microprocessor clock.

Final Progress Report

An accurate phase jitter simulation method was developed, which includes the phase jitter model in transient simulations. Employing current-steering logic, we designed, fabricated and tested a low noise PLL clock generator in a 0.5 μm CMOS process. This design, which benefited from availability of the jitter simulator, was also fabricated and tested. It achieves a top frequency of nearly 800 MHz with a power supply voltage of 1.8 V, a measured absolute phase jitter of less than 60 ps, and an RMS cycle-to-cycle phase jitter of 10 ps. This was the best phase jitter performance at that time, and it was achieved with low-voltage techniques which will have direct applicability to future CMOS circuits.

Several CAD tools were developed which support the design of advanced integrated circuits such as those from the PUMA project. A high-level optimization tool called GAIN (Genetic Algorithm on the INternet), was developed to assist a designer in judiciously allocating resources and partitioning logic onto chips in MCM designs. It uses a genetic algorithm to explore permutations of a baseline architecture, spawning trace-driven simulation jobs on a network of workstations so that many options can be evaluated in parallel.

Our subcontractor, Cascade Design Automation, developed a cell library migration tool called MasterPort, which converts a GDSII input layout to compacted layout in a specified rule set. The tool automatically generates the constraints and solves the constraint equations. It was used in the development of more than 120 cells for test chips designed in the project, and was very helpful in keeping the cells updated in the rapidly evolving CGaAs process. Cascade also developed an area-distributed pad router, called Eggo, which worked with existing placement tools to minimize power and signal routing between the array of bumps on the surface of a chip and the modules to which they are connected.

TEMPO is a transistor-level micro-placement tool for two-dimensional cell synthesis. It generates custom-quality layouts for such high-performance logic families as cascode voltage switch logic, pass transistor logic, and domino CMOS. This is achieved through powerful transformations such as dynamic geometry sharing through transistor chaining and arbitrary geometry merging. TEMPO enables the quick migration of cell libraries to new fabrication processes.

A constructive logic synthesis tool, called M31, was developed to interleave the traditionally separate technology-independent logic restructuring and technology-dependent library binding stages of circuit synthesis. M31 is based on Boolean decomposition strategy that ties together 1) the structural properties of the functions being synthesized, 2) the structural attributes of the implementation network, and 3) the functional content of the target library. The resulting implementations are consistently smaller and faster than those generated using conventional logic synthesis. In addition, they can be incrementally modified to create variants that achieve other area/speed trade-offs.

A methodology and tools for minimizing the effects of capacitively-coupled crosstalk were also developed. By using an accurate and consistent empirical model for wiring resources and constraints, coupled noise and delay were made predictable, and thus avoidable. A congestion-driven placement algorithm was developed to help minimize the incidence of capacitive coupling, and a global route-embedder was developed to guide the detailed router to meet timing and noise constraints.

Papers on each of these topics are included in the list of manuscripts attached. Presentations and project details can be found at <http://www.eecs.umich.edu/UMichMP/>.

Conclusions

Many of the characteristics of CGaAs make it an ideal technology for space-based applications. Unlike other GaAs technologies, it has a p-transistor, which facilitates efficient on-chip memory, and provides most of the other benefits of CMOS. Like other GaAs technologies, CGaAs is generations behind CMOS in scaling, which means that it cannot compete with CMOS for speed. The power-delay product of CGaAs, though, is extremely good compared to a similar generation of CMOS, and its radiation hardness is superb. CGaAs devices do have more gate and drain leakage than CMOS. In most respects, CGaAs scales well; there is no gate oxide to scale, the uniformity of which will be a serious challenge for CMOS below a certain thickness. On the other hand, making source and drain contact to the epitaxial material is difficult, so scaling this contact area is a challenge. And finally, the process control was not tight enough, and not well enough defined to yield fully functional circuits. Nevertheless, a number of useful contributions from this project in computer architecture, circuit design, packaging and CAD tools were generated in the PUMA project.

Final Progress Report

List of Publications

- "A Constraint-Driven Compiler for Process-Tolerant SRAMs," Richard B. Brown, Ajay Chandna, C. David Kibler, Mark Roberts, and Karem A. Sakallah, IEEE Transactions on Very Large Scale Integrated (VLSI) Systems, accepted for publication.
- "Performance Limits of Trace Caches," M. Postiff, G. Tyson, T. Mudge, The Journal of Instruction-Level Parallelism, accepted for publication.
- "Timing Verification of Sequential Domino Circuits," by David Van Campenhout, Trevor Mudge and Karem A. Sakallah, IEEE Transactions on Computer-Aided Design, 18(5), May, 1999, pp. 645-658.
- "Transistor Level Micro-Placement and Routing for Two-Dimensional Digital VLSI Cell Synthesis," M.A. Riepe, K.A. Sakallah, International Symposium on Physical Design, April 12-14, 1999, pp. 74-81.
- "Crosstalk Constrained Global Route Embedding," Phiroze Parakh and Richard B. Brown, International Symposium on Physical Design, April 12-14, 1999, pp. 201-206.
- "A Quantitative Approach to Non-linear Process Design Rule Scaling," Spencer M. Gold, Bruce Bernhardt, Richard B. Brown, Advanced Research in VLSI, March 21-24, 1999, pp. 99-112.
- "A Complementary GaAs Microprocessor for Space Applications," T. Basso, R. B. Brown, The Second International Conference on Integrated Micro-Nanotechnology for Space Applications, November 1998.
- "The Edge-Based Design-Rule Model Revisited," M. Riepe, K. Sakallah, ACM Trans. on Design Automation of Electronic Systems, Vol. 3, no. 3, July 1, 1998, pp. 463-486.
- "Congestion Driven Quadratic Placement," P. Parakh, R. B. Brown, K. Sakallah, 35th Design Automation Conference, San Francisco, CA, June 15-19, 1998, pp. 275-278.
- "M32: A Constructive Multilevel Logic Synthesis System," V. N. Kravets and K. A. Sakallah, 35th Design Automation Conference, June 15-19, 1998, pp. 336-341.
- "Overview of complementary GaAs technology for high-speed VLSI circuits," R. Brown, B. Bernhardt, M. LaMacchia, J. Abrokwhah, P. Parakh, T. Basso, S. Gold, S. Stetson, C. Gauthier, D. Foster, B. Crawford, T. McQuire, K. Sakallah, R. Lomax, T. Mudge, IEEE Transactions on VLSI Circuits, vol. 6, no. 1, pp. 47-51, March 1998.
- "AFTA: A Delay Model for Functional Timing Analysis," V. Chandramouli, J. P. Whittemore, K. A. Sakallah, Proc. Design Automation and Test in Europe, February 23-26, 1998. pp. 350-355.
- "Improving Code Density Using Compression Techniques," C. Lefurgy, P. Bird, I-C. Chen, and T. Mudge, Proceedings of the 29th Annual International Symposium on Microarchitecture, pp. 194-203, December 1997.
- "The bi-mode branch predictor," C. Lee, I. Chen, and T. Mudge, 29th Ann. IEEE/ACM Symp. Microarchitecture (MICRO-29), pp. 4-13, Dec. 1997.
- "A Complementary GaAs (CGaAs) (TM) 32-bit Multiply Accumulate Unit," M. J. Kelley, M. A. Postiff, T. D. Strong, R. B. Brown, and T. N. Mudge, 31st Asilomar Conference on Signals, Systems, and Computers, November 2-5, 1997.
- "Multilevel optimization of pipelined caches," O. Olukotun, T. Mudge, R. Brown, IEEE Trans. on Computers, vol. 46, no. 10, pp 1093-1101, Oct. 1997.

Final Progress Report

- "Instruction prefetching using branch prediction information," I-C. Chen, C-C. Lee, and T. Mudge, International Conference on Computer Design 97, pp. 593-601, Oct. 1997.
- "Design Optimization for High-Speed Per-Address Two-Level Branch Predictors," I-Cheng K. Chen, Chih-Chieh Lee, Matt Postiff, and Trevor Mudge, International Conference on Computer Design (ICCD), Austin, Texas, pp. 88-96, October 12-15, 1997.
- "Choosing the Appropriate Thresholds for Measuring Propagation Delay and Transition Time," V. Chandramouli, K. A. Sakallah, AICSP Special Issue, Analog Issues in Digital VLSI, E. B. Friedman, ed., Kluwer Academic Publishers, pp. 9-28, September 1997.
- "Selection of Voltage Thresholds for Delay Measurement," V. Chandramouli, Karem A. Sakallah, Analog Integrated Circuits and Signal Processing: An International Journal, Vol. 14, Number 1/2, September 1997, pp. 9-28.
- "Signal Delay in Coupled Distributed RC Lines in the Presence of Temporal Proximity," V. Chandramouli, A. I. Kayssi, and K. A. Sakallah, Proceedings of the Advanced Conference in VLSI, Ann Arbor, Michigan, pp. 32-46, September 15-16, 1997.
- "Improving data cache performance by pre-executing instructions under a cache miss," J. Dundas and T. Mudge, Proc. 1997 ACM Int. Conf. on Supercomputing, pp. 68-75, July 1997, pp. 68-75.
- "Impact of MCMs on high performance processors," B. Davis, C. Gauthier, P. Parakh, T. Basso, C. Lefurgy, R. Brown, and T. Mudge, Proc. ASME Advances in Electronic Packaging 97, vol. 1 (EEP-vol. 19-1), pp. 863-868, June 1997.
- "Trace-driven memory simulation: A survey," R. Uhlig and T. Mudge, ACM Computing Surveys, vol. 29, no. 2, pp. 128-170, June 1997.
- "A Discussion of a GaAs MCM Fabricated at MicroModule Systems through the Multichip Module Designer's Access Service (MIDAS)," J. Peltier, W. Hansford, C. Gauthier, R. Lomax, M. Nanua, P. Parakh, S. Stetson, Proceedings of 6th International Conference on Multichip Modules, Denver, CO, Apr. 2-4, 1997.
- "Computation of Switching Noise in Printed Circuit Boards," J. G. Yook, V. Chandramouli, L. P. Katehi, K. A. Sakallah, T. Arabi, T. Schreyer, IEEE Transactions on Components, Packaging, and Manufacturing Technology, vol. 20, no. 1, pp. 64-75, March 1997.
- "Software-managed address translation," B. L. Jacob and T. N. Mudge, Proceedings: Third International Symposium on High Performance Computer Architecture (HPCA-3), San Antonio Texas, February 1997.
- "Area I/O Flip-Chip Packaging to Minimize Interconnect Length," R. J. Lomax, R. B. Brown, M. Nanua, T. D. Strong, Proceedings of 1997 IEEE Multi-Chip Module Conference, Santa Cruz, CA, pp. 2-7, Feb. 1997.
- "CAD Tools for Area-Distributed I/O Pad Packaging," R. Farbarik, X. Liu, M. Rossman, P. Parakh, T. Basso, R. Brown, 1997 IEEE Multi-Chip Module Conference, February 4-5, 1997, pp. 125-129.
- "Trap-driven memory simulation with Tapeworm II," R. Uhlig, D. Nagle, T. Mudge and S. Sechrest, ACM Trans. Modeling and Computer Simulation (TOMACS), vol. 7, no. 1, pp. 7-41, Jan. 1997.
- "Access to Local Resources in a Nomadic Environment," B. Jacob and T. Mudge, USENIX Technical Conference, Anaheim, CA, January 6-10, 1997.

Final Progress Report

"Timing Verification of Sequential Domino Circuits," D. Van Campenhout, T. Mudge, K.A. Sakallah, ICCAD 96: IEEE/ACM Digest of Technical Papers, San Jose, CA, November 10-14, 1996, pp. 127-132.

"Complementary GaAs Technology for a GHz Microprocessor," R. Brown, T. Basso, P. Parakh, S. Gold, C. Gauthier, R. Lomax and T. Mudge, 18th Annual IEEE GaAs IC Symposium: Technical Digest 1996, Orlando, FL, Nov. 3-6, 1996, pp. 313-316.

"A Complementary GaAs PLL Clock Multiplier with Wide-Bandwidth and Low-Voltage Operation," P. Stetson and R. Brown, 18th Annual IEEE GaAs IC Symposium: Technical Digest 1996, Orlando, FL, Nov. 3-6, 1996, pp. 317-320.

"An Analytical Model for Designing Memory Hierarchies," B.L. Jacob, P.M. Chen, S.R. Silverman, T.N. Mudge, IEEE Transactions on Computers, vol. 45, no. 10, October 1996, pp. 1180-1194.

"Analysis of Branch Prediction via Data Compression," I-C. K. Chen, J.T. Coffey, T.N. Mudge, ASPLOS VII: Architectural Support for Programming Languages & Operating Systems, October 1996, pp. 128-137.

"Support for Nomadism in a Global Environment," B. Jacob, T. Mudge, The Eleventh Annual ACM Conference on Object-Oriented Programming Systems, Languages and Applications, San Jose, CA, October 6-10, 1996.

"Timing Verification of Sequential Domino Circuits," D. Van Campenhout, T. Mudge, K.A. Sakallah, Techcon '96, Phoenix, AZ, September 12-14, 1996, no page numbers.

"The trading function in action," B. Jacob, T. Mudge, Proceedings of the Seventh ACM SIGOPS European Workshop, Connemara, Ireland, September 8-11, 1996, pp. 241-247.

"Modeling the Effects of Temporal Proximity of Input Transitions on Gate Propagation Delay and Transition Time," V. Chandramouli, K.A. Sakallah, 33rd Design Automation Conference, Las Vegas, NV, June 1996, pp. 617-622.

"Correlation and Aliasing in Dynamic Branch Predictors," S. Sechrest, C-C Lee, T. Mudge, ISCA '96 Proceedings: 23rd Annual International Symposium on Computer Architecture, May 1996, pp. 22-32.

"A Comparison of Two Pipelined Structures," M. Golden, T. Mudge, IEEE Proceedings: Computers and Digital Techniques, vol. 143, no. 3, May 1996, pp. 161-167.

"Ravel-XL: A Hardware Accelerator for Assigned-Delay Compiled-Code Logic Gate Simulation," M.A. Riepe, J.P. Marques Silva, K.A. Sakallah, R.B. Brown, IEEE Transactions on Very Large Scale Integration (VLSI) Systems, vol. 4, no. 1, March 1996, pp. 113-129.

"The Role of Adaptivity in Two-Level Adaptive Branch Prediction," S. Sechrest, C.C. Lee, T. Mudge, Proceedings of the 28th Annual IEEE ACM Symposium on Microarchitecture, MICRO-28, Ann Arbor, MI, 11/29-12/1/95, pp. 264-269.

"Power Rail Logic: A Low Power Logic Style for Digital GaAs Circuits," A. Chandna, R. B. Brown, D. Putti, and C. D. Kibler, IEEE Journal of Solid-State Circuits, vol. 30, no. 10, Oct. 1995, pp. 1096-1100.

A Parallel Genetic Algorithm for Multiobjective Microprocessor Design, by T.J. Stanley and T. Mudge, 6th Int. Conf. on Genetic Algorithms, Pittsburgh, PA, July, 1995.

"Instruction Fetching: Coping with Code Bloat," R. Uhlig, D. Nagle, T. Mudge, S. Sechrest, J. Emer, The 22nd Annual International Symposium on Computer Architecture, June 22-24, 1995, pp. 345-356.

Final Progress Report

"The Aurora RAM Compiler," A. Chandna, C.D. Kibler, R.B. Brown, M. Roberts, K.A. Sakallah, 32nd Design Automation Conference: Proceedings 1995, San Francisco, CA, June 12-16, 1995 pp. 261-266.

"Critical Paths in Circuits with Level-Sensitive Latches," T. Burks, K.A. Sakallah, T.N. Mudge, IEEE Transactions on Very Large Scale Integration (VLSI) Systems, vol. 3, no. 2, June 1995, pp. 273-291.

"Timing Models for Gallium Arsenide Direct-Coupled FET Logic Circuits," A.I. Kayssi, K.A. Sakallah, IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, vol. 14, no. 3, March 1995, pp. 384-393.

"Systematic Objective-driven Computer Architecture Optimization," T. Stanley, T. Mudge, 16th Conference on Advanced Research in VLSI, Chapel Hill, NC, March 27-29, 1995, pp. 286-300.

"A Verilog Preprocessor for Representing Datapath Components," B.T. Davis, T. Mudge, 4th International Verilog HDL Conference: Proceedings 1995, March 27-29, 1995, pp. 90-98.

"An Asynchronous GaAs MESFET Static RAM Using a New Current Mirror Memory Cell," A. Chandna and R. Brown, IEEE Journal of Solid State Circuits, vol. 29, no. 10, pp. 1270-1276, Oct. 1994.

"Trap-Driven Simulation with Tapeworm II," R. Uhlig, D. Nagle, T. Mudge and S. Sechrest, 6th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-VI), pp.132-144, October 1994.

"Resource Allocation in a High Clock Rate Microprocessor," M. Upton, T. Huff, T. Mudge and R. Brown, 6th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-VI), pp. 98-109, October 1994.

"A Variable-Voltage Bidirectional I/O Pad for Digital GaAs Applications," P.J. Sherhart, M.D. Upton, R.L. Lomax, R.B. Brown, Proceedings of the 1994 IEEE GaAs IC Symposium, pp. 67-70, October 16-19, 1994.

"Power Rail Logic: A Low Power Logic Style for Digital GaAs Circuits," A. Chandna, R. B. Brown, D. Putti, C. D. Kibler, Proceedings of the 1994 IEEE GaAs Symposium, pp. 71-74, October 16-19, 1994.

Manuscripts previously submitted to ARO, paper not accepted

"Constructive Library-Aware Synthesis Using Symmetries," V.N. Kravets, K. Sakallah, ICCAD 1999.

"New Benchmark Circuits for Problems in Cell-Level Digital VLSI Design Automation," M. Riepe, K. Sakallah, submitted to 1999 International Symposium on Physical Design.

"Satisfiability-based FPGA Routing," G.J- Nam, K.A. Sakallah, R.A. Rutenbar, 1998 Int. Conference on Computer Aided Design.

"Transistor level microplacement and routing for two-dimensional digital VLSI cell synthesis," M. Riepe, K. Sakallah, 1998 Int. Conference on Computer Aided Design.

"A Satisfiability-Based Algorithm for Pseudo-Boolean Optimization with Applications to Layout Synthesis," M. Riepe, J. Silva, K. Sakallah, ISPD'97.

"A Source Code Study of Predicated Execution," M. Golden, T. Mudge, IEEE Trans. on Computers.

"Software-Managed Address Translation," T. Mudge, B. Jacob, ASPLOS 1996.

Final Progress Report

"Distributed Objects, Descriptive Service Lookup, and Dynamic Interfaces - The Building Blocks of an Extensible Wide-Area Transaction Environment," B. Jacob, ACM SOSP, 1995.

"Implementing IEEE Rounding in Parallel-Array Floating-Point Multipliers," M. Riepe, T. Huff, T. Mudge, 1995 Symp. on Computer Arithmetic.

Technical Reports

"A Quantitative Approach to Nonlinear IC Process Design Rule Scaling," S. M. Gold, Ph.D. dissertation, University of Michigan, 1999, Tech. Report, SSEL-288.

"Improving Processor Performance by Dynamically Pre-Processing the Instruction Stream," J. Dundas, Ph.D. dissertation, University of Michigan, 1999, Tech. Report, CSE-TR-389-99.

"Transistor Level Micro Placement and Routing for Two-Dimensional Digital VLSI Cell Synthesis," M. Riepe, Ph.D. dissertation, University of Michigan, 1999, Tech. Report, CSE-TR-390-99.

"A Microarchitecture for Resource-Limited Superscalar Microprocessors," T. Basso, Ph.D. dissertation, University of Michigan, 1999, Tech. Report SSEL-289.

"Design Considerations for Low Phase Jitter Clock Generators," P.S. Stetson, Ph.D. dissertation, University of Michigan, 1999, Tech. Report, SSEL-290.

"A Design Methodology for Addressing Crosstalk in Integrated Circuits," P. Parakh, Ph.D. dissertation, University of Michigan, 1999, Tech. Report, SSEL-291.

"A Design Methodology for Minimizing Crosstalk," P.N. Parakh, Ph.D. proposal, The University of Michigan, January 23, 1998.

"Improving Processor Performance by Dynamically Pre-processing the Instruction Stream," J.D. Dundas, Ph.D. proposal, The University of Michigan, 1997.

"High Speed Processor-Memory Interfaces," C. Gauthier, Ph.D. proposal, The University of Michigan, 1997.

"A Microarchitecture for High-Speed, Resource-limited, Superscalar Microprocessors," T.D. Basso, Ph.D. proposal, The University of Michigan, October 30, 1997.

"Transistor-level Placement and Routing Techniques for Digital VLSI Cell Synthesis," M.R. Riepe, Ph.D. proposal, The University of Michigan, September 17, 1997.

"The Design and Optimization of CGaAs Cache Memory," S.M. Gold, Ph.D. proposal, The University of Michigan, August 14, 1997.

"Low Jitter Clock Distribution Networks," S. Stetson, Ph.D. proposal, The University of Michigan, July 17, 1997.

"Architectural and Circuit Issues for a High Clock Rate Floating-Point Processor," T.R. Huff, Technical Report, SSEL-251, University of Michigan, 1995.

"GaAs MESFET Static RAM Design For Embedded Applications," A. Chandna, Technical Report, SSEL-252, University of Michigan, 1994.

Final Progress Report

List of Participants

Scientific Personnel Supported by this Project and Degrees Awarded:

<u>Faculty</u>	<u>e-mail address</u>	<u>Phone</u>	<u>FAX</u>
Richard B. Brown	brown@umich.edu	734-763-4207	734-763-9324
Ronald J. Lomax	rjl@engin.umich.edu	734-936-2972	734-747-1781
Trevor N. Mudge	trn@eecs.umich.edu	734-764-0203	734-763-4617
Karem A. Sakallah	karem@eecs.umich.edu	734-936-1350	734-763-4617

<u>Students</u>	<u>e-mail address</u>	<u>Degrees Granted During Project</u>	<u>Year</u>
Todd Basso	todd.basso@sun.com	M.S. Ph.D.	1996 1999
Jay Cameron	cameronj@eecs.umich.edu		
Ajay Chandna	ajayc@lakewood.sps.mot.com	Ph.D.	1995
V. Chandramouli	chandra@austin.ibm.com	Ph.D.	1998
I-Cheng Chen	kevin.chen@amd.com	Ph.D.	1997
Koushik Das	kdas@engin.umich.edu		
Brian Davis	btdavis@eecs.umich.edu	M.S.	1994
Alan Drake	ajdrake@engin.umich.edu		
James Dundas	dundas@eecs.umich.edu	M.S. Ph.D.	1995 1998
Krisztian Flautner	manowar@engin.umich.edu	M.S.	1998
Claude R. Gauthier	clauderg@engin.umich.edu	M.S. Ph.D.	1997 1999
Spencer Gold	spencer.gold@sun.com	M.S. Ph.D.	1997 1999
Michael Golden	michael.golden@amd.com	Ph.D.	1995
Matthew Guthaus	mguthaus@eecs.umich.edu	B.S.E.	1998
John Hall	hallj@engin.umich.edu	M.S.	1999
Rob Hower	hower@engin.umich.edu	Ph.D.	2000 ant.
Tom Huff	thuff@ichips.intel.com	Ph.D.	1995
Bruce Jacob	blj@beckmann.eng.umd.edu	M.S. Ph.D.	1995 1997
Michael Joseph Kelley	not available	M.S.	1997
Brian Kelly	not available	M.S.	1997
Keith Kraver	kkraiver@engin.umich.edu	M.S.	1997

Final Progress Report

Victor Kravets	vkavets@eecs.umich.edu	M.S.	1993
Chih-Chieh Lee	cclee@pa.dec.com	M.S.	1995
		Ph.D.	1998
Ting-Leung Lee	not available	M.S.	1996
Charles Lefurgy	lefurgy@eecs.umich.edu	M.S.	1996
Vince Mazzotta	vmazzott@tellabs.com	M.S.	1996
David Nagle	bassoon@ece.cmu.edu	Ph.D.	1995
Mini Nanua	nanuam@engin.umich.edu	M.S.	1996
Matthew A. Postiff	postiffm@eecs.umich.edu	M.S.	1997
Phiroze N. Parakh	parakh@mondes.com	M.S.	1996
		Ph.D.	1999
Michael A. Riepe	riepe@magma-da.com	Ph.D.	1999
Mark Roberts	Mark.Roberts@amd.com	M.S.	1994
Jose Robins	not available		
Arvind Salian	arvi@engin.umich.edu		
Himanshu Sharma	hsharma@engin.umich.edu	M.S.	1999
Patrick Sherhart	pjsherhart@macconnect.com	M.S.	1994
P. Sean Stetson	sstetson@ti.com	M.S.	1996
		Ph.D.	1998
Tim Strong	strongtd@engin.umich.edu	M.S.	1997
Michael D. Upton	mupton@ichips.intel.com	Ph.D.	1997
David Van Campenhout	davidvc@eecs.umich.edu	Ph.D.	1999
Anand Varadharajan	anandv@pa.dec.com	M.S.	1997
John Wei	johnwei@umich.edu	M.S.	1999

Final Progress Report

Report of Inventions

Ajay Chandna and Richard B. Brown, U.S. Patent 5,490,105, "High Speed Current Mirror Memory Cell Architecture," Feb. 6, 1996.

Constructive Multilevel Logic Synthesis, disclosed to UM Technology Transfer Office.

Active Current Mirror-Based Current-Mode Transceivers, to be disclosed to UM Technology Transfer Office.

Bibliography

1. Linley Gwennap, "The Death of the Superprocessor," Microprocessor Report, vol. 9, no. 13, p. 3, October 2, 1995.
2. Michael D. Upton, Architectural Trade-offs in a Latency Tolerant Gallium Arsenide Microprocessor, Ph.D. Dissertation, University of Michigan, Ann Arbor, 1997.
3. B. Bernhardt, M. LaMacchia, J. Abrokwhah, J. Hallmark, R. Lucero, B. Mathes, B. Crawforth, D. Foster, K. Clauss, S. Emmert, T. Lien, E. Lopez, V. Mazzotta, B. Oh, "Complementary GaAs (CGaAsTM): A High Performance BiCMOS Alternative," GaAs IC Symposium Technical Digest, San Diego, CA, Oct. 29–Nov. 1, 1995, pp. 18–21.
4. J. Abrokwhah, J. Huang, W. Ooms, C. Shurboff, J. Hallmark, R. Lucero, J. Gilbert, B. Bernhardt, and G. Hansell, "A Manufacturable Complementary GaAs Process," GaAs IC Symposium Technical Digest, San Jose, CA, Oct. 10–13, 1993, pp. 127–130.
5. P. O'Neil, B. Bernhardt, F. Nikpourian, C. Della, Y. Abad, G. Hansell, "GaAs Integrated Circuit Fabrication at Motorola," GaAs IC Symposium Technical Digest, San Jose, CA, Oct. 10–13, 1993, pp. 123–126.
6. J. Hallmark, C. Shurboff, W. Ooms, R. Lucero, J. Abrokwhah, and J. Huang, "0.9V DSP Blocks, A 15ns 4K SRAM and a 45ns 16-bit Multiply/Accumulator," GaAs IC Symposium Technical Digest, 1994, pp. 55–58.

REPORT DOCUMENTATION PAGE			<i>Form Approved</i> OMB NO. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comment regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE December 19, 1999		3. REPORT TYPE AND DATES COVERED 27 September 1994-27 September 1999
4. TITLE AND SUBTITLE Design Optimization of a GaAs RISC Microprocessor with Area-Interconnect MCM Packaging			5. FUNDING NUMBERS DAAH04-94-G-0327	
6. AUTHOR(S) Richard B. Brown				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Michigan 1301 Beal Ave. Ann Arbor, MI 48109-2122			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211			10. SPONSORING / MONITORING AGENCY REPORT NUMBER ARO 33 790.78-EL	
11. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution unlimited.			12 b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) This project analyzed Complementary Gallium Arsenide (CGaAs) and advanced packaging technologies for use in high performance and radiation hard circuits. The basic CGaAs process was analyzed in this project for non-linear design rule scaling. Static, domino and dual-rail domino (CVSL) circuits were designed to evaluate CGaAs for use in VLSI circuits. Phase-Locked Loop and current-mode I/O circuits were designed and tested. To facilitate the design of systems proposed in this project, a CGaAs cell library, SRAM compiler, and place-and-route tools that support flip-chip area I/O packaging were developed. A gold-bumping process was developed in the UM solid-state electronics laboratory which produces bumps on pitches as tight as 50 μ m. A superscalar PowerPC microarchitecture was developed for implementation in the modest integration levels of CGaAs. The project culminated in the design and testing of the PUMA PowerPC integer processor, which incorporates area-I/O for flip-chip packaging. Parameter variation in the CGaAs process of the prototype run rendered the unipolar-logic decoder circuits in the SRAMs inoperative; nevertheless, most of the processor was functional. This project demonstrated that CGaAs is a viable technology for radiation-hard microprocessors, but it would need to have threshold voltages and minimum geometries scaled to achieve high performance.				
14. SUBJECT TERMS Radiation-Hard, Gallium Arsenide, Microprocessor			15. NUMBER OF PAGES	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OR REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	